REMARKS

Claims 1 to 3, 5 to 41 and 43 to 46 are pending after this amendment.
Claims 4 and 42 are cancelled.

Claim 1 has been amended to explicitly recite that the compute nodes are operable to execute application processes and to more specifically describe encapsulation and address translation. Similar amendments have been made to other independent claims. These amendments are submitted to not narrow the scope of the claims. The text added to the claims describing compute nodes, address translation and encapsulation is submitted to be inherent according to the ordinary meanings of compute node, encapsulation, and address translation in the field of this invention. Other clarifying amendments have been made to the claims.

The Applicants' agent respectfully requests a telephone interview with the Examiner and the Examiner's supervisor for the purpose of discussing the distinctions of the claimed technology over the cited references and for better understanding the positions taken by the Examiner. The Examiner is asked to contact the Applicants' agent to arrange a time and date for the interview.

Compliance with 35 U.S.C. §102

The Office Action alleges that claims 1-11 and 19-46 are anticipated by Pettey (US 7046668). The pending claims are submitted to patentably distinguish Pettey at least for the reasons set out below.

Claims 1 to 11 and 19 to 38

Claim 1 recites a method that involves transport of a local packetized interconnect packet from the local packetized interconnect of a sending compute node to the local packetized interconnect of a receiving compute node. The local packetized interconnect packet is encapsulated in a data payload of an inter-node communication network packet as it is transmitted between an interface of the sending compute node to an interface of the receiving compute node. Claims 2-11 and 19 to 38 all depend from independent claim 1.

The Examiner is encouraged to review paragraphs [0029] to [0031] of the present application which describe a motivation for the claimed invention. Applicant submits that Pettey does not address the latency issues which motivate the claimed invention.

For Pettey to anticipate claim 1, Pettey would need to provide the following sequence of steps:

- placing a local packetized interconnect packet on the local packetized interconnect of the sending compute node;
- receiving the local packetized interconnect packet at the network interface of the sending compute node;
- encapsulating the local packetized interconnect packet in an inter-node communication network packet addressed to the receiving compute node by placing the local packetized interconnect packet in a data payload of the inter-node communication network packet;
- dispatching the inter-node communication network packet to the receiving compute node by way of the inter-node communication network;
- receiving the inter-node communication network packet at the network interface of the receiving compute node;
- extracting the local packetized interconnect packet from the inter-node communication network packet; and,
- placing the extracted packet onto the local packetized interconnect of the receiving compute node.

As discussed in detail below, Pettey does not disclose the claimed sequence of steps.

Pettey

Pettey describes some systems in which servers each have dedicated I/O devices (Figs. 1 to 3) and some systems in which several servers can share the same I/O devices by way of a shared I/O switch (Figs. 4 to 8, 13, 14, 19 and 20).

Note that Pettey, as understood, does not provide for the exchange of data between any two servers connected to the same shared I/O switch. Pettey uses the OS header mechanism to keep packets associated with their respective originating root complexes (col. 4, line 30 - col. 5, line 47; col. 6, lines 4 - 35). At the shared I/O endpoint, separate control registers are provided for each root complex (col. 5, lines 48 - 59) and separate resources are available to each OS Domain (col. 5, line 60 - col. 6, line 3). Thus, packets associated with each root complex are distinguished and kept separate from packets associated with other root complexes.

The following drawing, which is based upon Pettey Fig. 6, is an example of a system having a shared I/O switch. An Ethernet cloud and ‘other computer or device’ have been added.

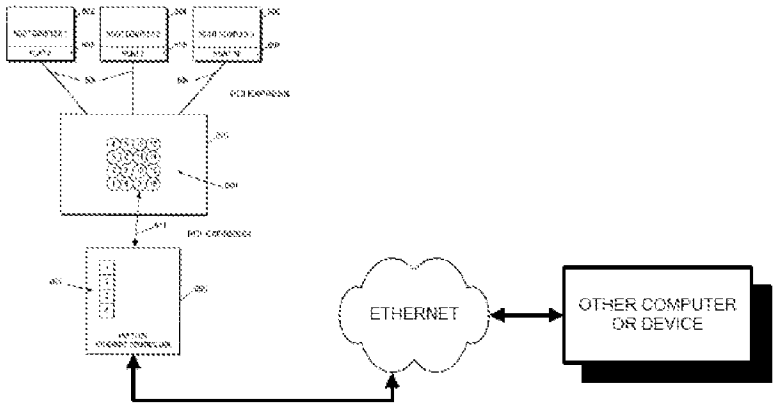

Pettey, as understood, describes communication between a processing complex upstream from one of the root complexes and some other computer or device connected to a network. Data can travel between the processing complex and the other computer or device by the path: processing complex ↔ root complex ↔ shared I/O switch ↔ shared Ethernet controller ↔ Ethernet ↔ other computer or device. In embodiments where the processing complex and root complex are not 'shared I/O aware', communication on links of this path can be:

- in a host bus protocol upstream from the root complex;
- in the PCI Express protocol between the root complex and shared I/O switch;
- in the PCI Express+ protocol between the shared I/O switch and the shared Ethernet controller; and
- in the Ethernet protocol between the Ethernet controller and the other computer or device.

Response to Office Action

The Applicant submits that the rejection of claim 1 in the Office Action errs because:

- it attempts to equate features shown by Pettey with claim elements in a way that is inconsistent with claim 1; and
- it equates different features of Pettey to the same claimed feature at different parts of the analysis.

Analysis of claim 1 requires identification of sending and receiving compute nodes.

The Office Action, as understood, begins by identifying the plurality of compute nodes as the servers (102, 104 and 106) shown in Fig. 1 of Pettey (Par. 4 of the Office Action refers to Fig. 1 and cites col. 7, ln. 54-67 which describes servers 102, 104 and 106). Fig. 1 of Pettey, as understood, shows three servers 102, 104 and 106 connected by Ethernet, Fiber Channel or Other networks. The Applicant points out that these servers cannot be equated to the claimed compute nodes because Pettey, as understood, does not indicate that servers 102, 104 and 106 tunnel data from a sending one of the servers to a receiving one of the servers as claimed in claim 1. For example, Pettey does not disclose or suggest that the embodiment shown in Figure 1 operates by encapsulating a local packetized interconnect packet at one server and extracting the local packetized interconnect packet at another server.

Later on in the analysis of claim 1, the Office Action appears to switch to equating one of the compute nodes to the shared I/O endpoint (on p. 3, the Office Action states that receiving the inter-node communication network packet at the network interface of the receiving compute node is disclosed at col. 18, ln. 64 to col. 19, ln 36). The cited section discusses the transmission of a PCI Express+ packets in either direction between a shared I/O switch and the shared I/O endpoint.

The shared I/O endpoint cannot properly be equated to the claimed sending or receiving compute node because it lacks an independent address space. Pettey specifically requires that "the shared I/O endpoint is considered to be within the load/store fabric of the OS Domain" (see col. 23, ln. 44-45 and col. 21, ln. 23-27). Pettey states that an OS Domain is an operating system domain wherein the system memory and I/O devices for a particular CPU (or set of CPU's) are part of a single system memory map or operating system (see col. 13, ln. 24-26). Pettey's load/store fabric, as understood establishes a single address space. Addresses within this address space are carried within the Header field of PCI Express and PCI Express+ packets (see Fig. 9 and 10). Pettey, as understood does not disclose or suggest changing any value in the Header field of PCI Express and PCI Express+ packets thus indicating that the packets remain within a single address space.

Further, claim 1 recites that the compute nodes are operable to execute application processes. The shared I/O endpoint does not provide this function. Shared I/O endpoints have specific dedicated

functions (Ethernet controller, Fiber Channel controller, other controller). Pettey's I/O controllers are not operable to execute application processes, as understood.

Further, claim 1 recites that the compute nodes each have their own local packetized interconnects. Pettey's shared I/O controller couples to a shared load/store fabric (1302 - See Fig. 13) belonging to the root complexes that share the shared I/O controller. Shared load/store fabric (1302) is external to the shared I/O controller (1300) and so cannot be a local packetized interconnect of the shared I/O controller. Shared I/O controller (1300) has a 'data path mux+' (1306). However, Pettey, as understood, does not disclose or suggest that the 'data path mux+' is a local packetized interconnect.

The Examiner alleges that Pettey discloses tunneling data from a sending to a receiving compute node, citing col. 23, ln. 35-40. This is submitted to be incorrect because:

- Claim 1 recites encapsulating the local packetized interconnect packet in an inter-node communication network packet at a network interface of a sending node and then receiving the inter-node communication network packet (with the encapsulated local packetized interconnect packet) at a network interface of the receiving node. Whatever tunneling Pettey provides is between structures that cannot properly be equated to these network interfaces.

As pointed out in a previous response, Pettey uses the words 'encapsulated' and 'tunneled' only in the following context:

> "It is also envisioned that the addition of a header within a load/store fabric, as described above, could be encapsulated within another load/store fabric yet to be developed, or could be encapsulated tunneled or embedded within a channel based fabric such as Advanced Switching or All Ethernet. Regardless of the fabric used downstream from the OS Domain (or root complex), the inventors consider any utilization of the method of associating a shared I/O endpoint with an OS Domain to be within the context of their invention, as long as the shared I/O endpoint is considered to be within the load/store fabric of the OS Domain" (see col. 23, ln. 35-45)

This section of Pettey, as understood, is specific to the communication of data between the root complex and a shared I/O endpoint.

Claim 1 recites that a local packetized interconnect packet is encapsulated in the payload of an inter-node communication packet at the network interface of a sending compute node. The inter-node communication packet must be received at the network interface of the receiving compute node.

Even if the above section of Pettey discloses encapsulating a local packetized interconnect packet in the payload of an inter-node communication packet (which is denied), the encapsulation occurs at the wrong location to match claim 1. Whatever encapsulation is provided by Pettey, the encapsulation must end at or upstream from the shared I/O endpoint (since the shared I/O endpoint uses the OS Domain Header).

The Office Action cannot equate the shared I/O endpoint to the receiving compute node because the Office Action, as understood, equates the shared I/O endpoint to the network interface of the sending compute node. Also, as noted above, the shared I/O endpoint is not a compute node.

The Office Action begins the claim 1 analysis by quoting col. 8., ln. 10 identifying the network interfaces of the claimed compute nodes with NIC (110) (see Office Action p. 3, ln. 1). However, later in the analysis of claim 1 the Office Action indicates that Pettey discloses the claimed step of 'receiving the local packetized interconnect packet at the network interface of the sending compute node' at col. 18, ln. 64-col. 19, ln. 11. This section of Pettey describes building a PCI Express+ packet and sending the PCI Express+ packet to the endpoint device such as a shared I/O Ethernet controller (see col. 19, ln. 4). In the cited passage, the only place that any packet is being received is at the shared I/O endpoint. Therefore, the Office Action is interpreted as equating the shared I/O endpoint to the network interface of the sending node.

Thus, whereas claim 1 recites that the local packetized interconnect packet is encapsulated in the inter-node communication packet after it is received at the network interface of the sending compute node and before the inter-node communication packet is dispatched, any tunneling or encapsulation described by Pettey at col. 23, ln. 35-40 occurs before a PCI Express+ packet being

sent by a server reaches the shared I/O Ethernet controller (which the Office Action equates to the claimed network interface).

The Office Action also alleges that Pettey discloses encapsulating the local packetized interconnect packet at col. 18, ln 64-col. 19, ln. 36. The cited section describes building a PCI Express + packet which includes an OS header. Pettey contrasts a PCI Express Packet (shown in Fig. 9) to a PCI Express + packet in (shown in Fig 10). Pettey's Figs. 9 and 10 are reproduced here for convenience:

Fig. 9

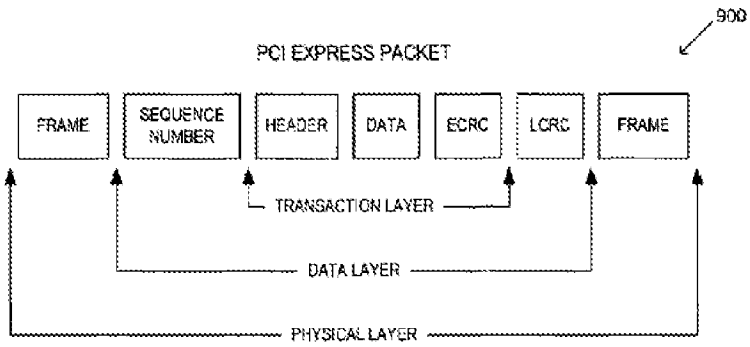

Fig. 10

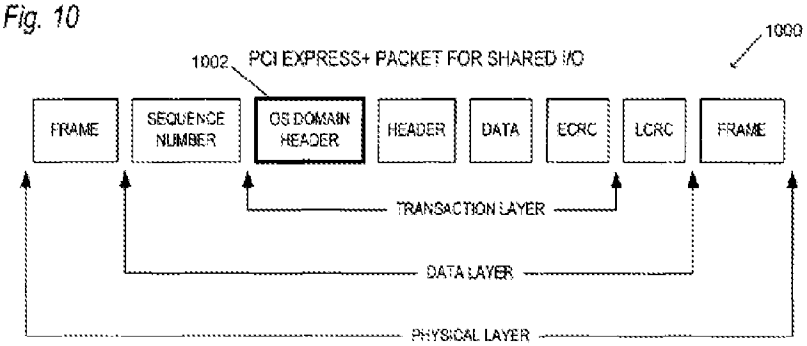

It can be seen that the difference between the PCI Express and PCI Express+ packets is the addition of an OS domain header (1002). Adding an OS domain header is not 'encapsulation'. Claim 1, as amended, specifically recites that encapsulation involves "placing the local packetized interconnect packet in a data payload of the inter-node communication network packet".

Further, the building of Pettey's PCI Express+ packets occur at the shared I/O switch (between the root complexes and shared I/O controllers). If the shared I/O controller is equated to the network interface of the sending compute node as alleged in the Office Action, as understood,

then building the PCI Express+ packet cannot be equated to the claimed encapsulating step because the PCI Express+ packet is built at the shared I/O switch, before the PCI Express+ packet reaches the shared I/O controller.

Other areas where the analysis of claim 1in the Office Action is submitted to be incorrect are:

- The Examiner indicates that Pettey shows compute nodes interconnected by PCI Express (Office Action par. 4, ln. 3). This is incorrect. Fig. 3 of Pettey, which the Examiner references, shows a single multi-CPU server (300) (see also the description of Fig. 3 at col. 11, ln. 27-42). Server (300) has internal PCI Express buses (320). Note that the solid lines joining servers (102), (104) and (106) to switches (122), (124) and (126) in Pettey Fig. 1 represent links operating under protocols of the ethernet, fiber channel or other network (these links are all downstream from I/O controllers (110), (112) and (114)). Therefore, in the embodiments of Figures 1 to 3, interconnections between Pettey's servers are by way of ethernet, fiber channel or other network links.
- It is incorrect for the Examiner to rely on features of Pettey's Figs 1 to 3 in combination with features associated with Pettey's shared I/O switches, shared I/O endpoints and added header. The embodiments of Pettey's Figs. 1 to 3, as understood, provide I/O interfaces for each server. These embodiments do not include shared I/O switches, shared I/O endpoints or conversion of PCI Express to PCI Express+ packets.
- The Office Action alleges that Pettey discloses 'dispatching the packet on the inter-node communication network' at col. 18, ln. 64 to col. 19, ln. 11. This section of Pettey describes building a PCI Express + packet and sending the PCI Express + packet to the endpoint device. However, the Office Action apparently equates the endpoint device (e.g. shared I/O controller) to the network interface of the sending compute node. Therefore, equating 'dispatching the packet on the inter-node communication network' to the description at col. 18, ln. 64 to col. 19, ln. 11 does not align with claim 1 because it would have the inter-node communication packet being dispatched from the network interface of the sending compute node before the arrival of the local packetized interconnect packet at the network interface.

For at least the reasons above, the Applicant submits that Pettey fails to anticipate claim 1. As all of claims 2 to 3 and 5 to 38 depend from claim 1, these claims are also submitted to patentably distinguish Pettey. The dependent claims are submitted to further distinguish Pettey for at least the reasons below.

Claim 2

Claim 2 recites providing an association at the network interface of a sending compute node. The association is between a first range of addresses (which are in an address space of the sending compute node) and a particular receiving compute node. Claim 2 also recites identifying the receiving compute node to which to address the inter-node communication packet in response to determining that the local packetized interconnect packet is addressed to the address in the first range of addresses.

The Office Action alleges that Pettey discloses, at a network interface, associating a first range of addresses with a receiving compute node. The Office Action alleges that this is shown at col. 16, ln. 7-31 and col. 18, ln 64 to col. 19, ln. 36. This is submitted to be incorrect. Pettey, as understood, discloses that a shared Ethernet controller (1300) can have a machine address for each OS domain it will be mapped to. This is an address of the shared ethernet controller. Pettey also says that the shared ethernet controller can have different sets of control registers corresponding to the different OS domains. The machine address of Pettey's ethernet controller, as understood, does not determine where the ethernet controller will send data. Apparently any PCI Express+ packet from an OS domain will be addressed to the same machine address regardless of the ultimate destination for the data in the PCI Express+ packet. Therefore, Pettey, as understood, fails to disclose the feature of claim 2.

The Office Action also alleges (see page 17) that Pettey's 'tables associated with domain' at col. 16, ln. 7-31 provide the first range of addresses claimed in claim 2. This portion of Pettey refers to a 'table which associates an OS Domain with a particular one of N number of operating system domain resources supported by the shared I/O controller'. Even if the contents of such a table are addresses, which is denied, the table identifies operating system domain resources supported by the shared I/O controller (apparently aspects of the controller itself). As noted above, the shared I/O controller cannot be equated to the receiving compute node. Operating system domain resources supported by the shared I/O controller cannot be equated to the receiving compute node either. Further, Pettey, as understood, does not disclose identifying a receiving compute node to which to address any inter-node communication packet in response to determining that the local packetized interconnect packet is addressed to the address in the first range of addresses.

Further, claim 2, as amended, recites "identifying the receiving compute node to which to address the inter-node communication packet in response to determining that the local packetized interconnect packet is associated with an address in the first range of addresses". Pettey, as understood, does not disclose this feature.

Therefore, claim 2 is submitted to be further patentable over Pettey.

Claim 3

Claim 3 recites "performing an address translation on the local packetized interconnect packet after receiving the packet at the network interface of the sending compute node and prior to placing the extracted packet onto the local packetized interconnect of the receiving compute node".

Pettey, as understood, does not disclose address translation. The Office Action appears to equate conversions between PCI Express packets and PCI Express + packets with the claimed address translation. This is incorrect at least because:

- It has been established above that the PCI Express and PCI Express+ links of Pettey belong to a single load/store fabric (col. 23, lines 35 - 45). A load/store fabric establishes a single address space. Addresses within this address space are carried within the Header field (Fig. 9 and 10) of PCI Express and PCI Express+ packets. Pettey does not disclose or suggest changing any values in the Header field.
- The claim recites that the address translation occurs 'after receiving the packet at the network interface of the sending compute node and prior to placing the extracted packet onto the local packetized interconnect of the receiving compute node'. By contrast, PCI Express + packets are created at the shared I/O switch before the packets reach the shared I/O controller that the Office Action has equated to the network interface of the sending compute node.
- Adding an OS domain header to a PCI Express packet is not address translation. The OS header identifies where (which OS Domain) the packet has come from, unlike an address which determines where a packet should go to.

Pettey specifically calls out that the load/store fabric of the root complex is extended to the shared I/O endpoint. For example, see:

- col. 13, ln. 22-26 (An OS Domain ... is an operating system domain where the system memory and I/O devices ... are part of a single system memory map ...)
- col. 3, ln. 53-57 (the solution should use a load/store architecture, where the processing complex sends data directly to ... or receives data directly from an I/O device ... without message passing ...)
- col. 23, ln. 41 to 45 ([T]he inventors consider any utilization of the method of associating a shared I/O endpoint with an OS Domain to be within the context of their invention, as long as the shared I/O endpoint is considered to be within the load/store fabric of the OS Domain)

Further, Claim 3, as amended recites that the address translation comprises "replacing an address of the local packetized interconnect packet in the address space of the sending compute node with a corresponding address in the address space of the receiving compute node". Pettey, as understood, does not disclose or suggest such an address translation.

Therefore claim 3 is submitted to be further patentable over Pettey.

Claim 5

Claim 5 recites that the address translation is performed at the network interface of the sending compute node. Even if converting between PCI Express and PCI Express+ were address translation (which is denied), in the Pettey system, conversion between PCI Express and PCI Express + occurs at a shared I/O switch. The Office Action, as understood, has equated the shared I/O controller, and not the shared I/O switch to the claimed network interface.

Therefore claim 5 is submitted to be further patentable over Pettey.

Claim 6

Claim 6 recites 'allocating a region of the memory of the receiving compute node to receive data from the sending compute node, memory locations in the allocated region being addressable by addresses in a second range of addresses corresponding to the first range of addresses, and communicating the second range of addresses from the receiving compute node to the sending compute node prior to tunneling the data from the sending compute node to the receiving compute node'.

The Office Action indicates that Pettey discloses these features as PCI Configuration logic and col. 16, ln. 7-31 and col. 18, ln. 64 to col. 19, ln. 36. This is submitted to be incorrect. The Applicants' representatives have carefully reviewed col. 18, ln. 64 to col. 19, ln. 36 and can find no disclosure relevant to the subject matter of claim 6. The reference to PCI configuration logic also does not disclose the feature of claim 6. PCI Configuration logic is generally understood as being logic that provides addresses of PCI components in one address space. PCI configuration logic, as understood, does not map between compute nodes. Pettey col. 16 ln. 7-31, as understood, describe how PCI Configuration logic can provide an address for a shared I/O controller for each of multiple root complexes.

Col. 16, ln. 7-31 indicates that the PCI Configuration logic controls the association of resources in an I/O controller to OS Domains and that the I/O controller may provide 'a hard-coded machine address' to the shared I/O switch for each OS Domain. PCI configuration logic, as understood, operates to assign addresses to PCI components in the address space of a PCI master (here the OS Domain). The machine address, as understood, is in the address space of the OS Domain. The cited section of Pettey does not describe 'communicating the second range of addresses from the receiving compute node to the sending compute node' as claimed.

Further, as noted above, Pettey's shared I/O controller is not a compute node and does not have a separate address space.

Therefore claim 6 is submitted to be further patentable over Pettey.

Claim 7

Claim 7 recites 'computing an address transformation between the first and second address ranges and using the address transformation to perform the address translation'. The Office Action alleges that this feature is disclosed at col. 16, ln. 7-31 and col. 18, ln. 64 to col. 19, ln. 36. The Applicants' representatives have carefully reviewed these sections of Pettey and can find no mention of computing an address transformation, as claimed.

Pettey's assignment of an OS Domain header to a packet being sent depends upon which OS Domain the packet is received from at the shared I/O switch. The OS Domain header is not derived by translating any address of the packet. In Pettey, as understood, the OS Domain number

is derived solely from the root complex associated with a packet. Pettey, as understood, does not disclose or suggest any mechanism by which the address of a PCI Express packet is used to derive the OS Domain header. Once an OS Domain has been assigned to a packet the shared I/O controller can use the OS Domain header to determine what resources are available for handling the packet.

Therefore claim 7 is submitted to be further patentable over Pettey.

Claim 10

Claim 10 recites that "the address translation is performed at the network interface of the receiving compute node". The Office Action alleges that this feature is disclosed by Pettey at col. col. 16, ln. 7-31 and col. 18, ln. 64 to col. 19, ln. 36. The Applicants assume that the Office Action is referring to Pettey's conversion between PCI Express and PCI Express+ which occurs at the shared I/O switch and involves adding or removing an OS Header. As noted above in relation to claim 3, conversion between PCI Express and PCI Express+ is submitted not to be address translation.

Further, as noted above, the Office Action, as understood, does not equate the shared I/O switch to a network interface.

Further, for packets destined from the shared I/O switch to an OS Domain, what is done is to strip off the OS Domain header (see e.g. col. 19, ln. 27; col. 20, ln. 65; and col. 21, ln. 19). Stripping off a header cannot properly be called performing an address translation between the first and second address ranges. Further, there is no need to compute an address translation if all one needs to do is to strip off a header.

Therefore claim 10 is submitted to be further patentable over Pettey.

Claim 11

Claim 11 has been amended to provide antecedent reference for 'address'. Claim 11, as amended recites "data to be written to a memory location in the memory system of the receiving compute node corresponding to an address of the local packetized interconnect packet in the address space of the sending compute node".

The Office Action alleged that the feature of claim 11 is disclosed by Pettey at col. 18, ln. 64 to col. 19, ln. 36. The Applicants' representatives have carefully reviewed the cited portion of Pettey and can find no reference to writing data to a memory location in the memory system of the receiving compute node corresponding to an address of the local packetized interconnect packet, as claimed.

Therefore claim 11 is submitted to be further patentable over Pettey.

Claims 19 to 38

The Examiner is asked to review claims 19 to 38 carefully. The Applicants' representatives have studied the reasons tendered in the Office Action for rejecting claims 19 to 38 but cannot understand how Pettey provides a factual basis for such claim rejections. Claims 19 to 38 are submitted to be patentable over Pettey both because they depend from claim 1 and because they recite features that further distinguish Pettey as discussed below.

Claim 19

Claim 19 recites that the local packetized interconnect packet comprises a read request packet. Contrary to the statement in the Office Action, nothing in col. 16, ln. 7-31 discloses a read request packet as claimed.

Therefore claim 19 is submitted to be further patentable over Pettey.

Claim 20

Claim 20 has been amended for clarity. Pettey is submitted not to disclose or suggest associating the first range of addresses with each of a plurality of receiving compute nodes as claimed in claim 20. The cited portion of Pettey (col. 7, ln. 9-31) discusses the assignment of a machine address to each OS Domain. Pettey, as understood, does not disclose or suggest assigning the same range of addresses at a source node to multiple receiving compute nodes, as claimed.

Therefore claim 20 is submitted to be further patentable over Pettey.

Claim 23

Claim 23 recites dispatching the inter-node communication network packet to the plurality of receiving compute nodes by way of a multicast feature of the inter-node communication network. The Office Action alleges that this feature is disclosed by Pettey at col. 6, ln. 5-24 (Ethernet). This is incorrect. The Examiner has already equated the inter-node communication network with PCI Express (Office Action paragraph 4). The Examiner cannot switch to equating the inter-node communication network with Ethernet for the purpose of claim 23. This is especially the case because Pettey, as understood does not disclose or suggest encapsulating any PCI Express or PCI Express+ packets (or packets of any other load/store fabric) in an ethernet or other packet.

Furthermore, The Applicants' representative can find no reference in Pettey to the use of a multicast feature on any network. The Office Action indicates that Pettey's Fig. 2 shows multicast however the Applicants' representatives can find no indication that this is the case.

Therefore claim 23 is submitted to be further patentable over Pettey.

Claim 24

Claim 24 recites "at the sending compute node, encapsulating each of a plurality of copies of the local packetized interconnect packet in a corresponding inter-node communication network packet addressed to a different one of the plurality of receiving compute nodes and dispatching each of the inter-node communication network packets to the corresponding receiving compute node by way of the inter-node communication network". The Office Action asserts that this is disclosed at col. 6, ln. 7-31. The Applicants' representatives can find no reference in this or any other section of Pettey to encapsulating each of a plurality of copies of the same local packetized interconnect packet in a corresponding inter-node communication network packet addressed to a different receiving compute node, as claimed.

Therefore claim 24 is submitted to be further patentable over Pettey.

Claim 25

Claim 25 recites "encapsulating the plurality of local packetized interconnect packets in the same inter-node communication network packet". The Office Action asserts that this is disclosed at col. 6, ln. 7-31. The Applicants' representatives can find no reference in this or any other section of

Pettey to encapsulating multiple local packetized interconnect packets in the same inter-node communication network packet.

Therefore claim 25 is submitted to be further patentable over Pettey.

Claim 26

Claim 26 recites "placing into the inter-node communication network packet the payloads of all of the plurality of local packetized interconnect packets and the header information from fewer than all of the plurality of local packetized interconnect packets". The Office Action asserts that this is disclosed at col. 6, ln. 7-31. The Applicants' representatives can find no reference in this or any other section of Pettey to placing into the inter-node communication network packet the payloads of all of the plurality of local packetized interconnect packets and the header information from fewer than all of the plurality of local packetized interconnect packets, as claimed.

Therefore claim 26 is submitted to be further patentable over Pettey.

Claims 27 and 28

Claim 27 recites that the sending and receiving compute nodes are peers. Claim 28 recites that the sending compute node has substantially the same construction as the receiving compute node. In the analysis offered in the Office Action, to the extent that it can be understood, the claimed sending compute node is equated with a server and the claimed receiving compute node is equated with the shared I/O controller or vice versa. The shared I/O controller and a server cannot be considered to be peers or to have substantially the same construction.

Therefore claims 27 and 28 are submitted to be further patentable over Pettey.

Claim 29

Claim 29 recites that the local packetized interconnect packet comprises an atomic read-modify-write packet. The Office Action alleges that this is provided by a PCI Express+ packet (col. 14, ln 32). Disclosure of a PCI Express+ packet is not disclosure of an atomic read-modify-write packet. A definition of read-modify-write packets is: "read-modify-write is a class of atomic operations such as test-and-set, fetch-and-add, and compare-and-swap which both read a memory location and write a new value into it simultaneously, either with a completely new

value or some function of the previous value." Pettey, as understood, does not discuss atomic read-modify-write packets either at col. 14, ln. 32 or elsewhere.

Therefore claim 29 is submitted to be further patentable over Pettey.

Claim 32

Claim 32 recites altering a correspondence of the first range of addresses to addresses in the address space of the receiving compute node by changing the second range of addresses to a third range of addresses. The Office action alleges that this feature is provided in Pettey at col. 16, ln. 7-31. The Applicants' representatives have carefully reviewed Pettey and do not see any reference in this or another section of Pettey to, having set up a correspondence between a first range of addresses at a sending compute node and a second range of addresses at a receiving compute node, changing that correspondence by changing the second range of addresses to a third range of addresses.

Therefore claim 32 is submitted to be further patentable over Pettey.

Claim 33

Claim 33 recites that the packet placed on the local packetized interconnect of the sending compute node is a read response packet and the method comprises converting the read response packet into a write request packet. The Office Action alleges that Pettey discloses this feature at col. 18, ln 55 - col. 19, ln. 36. The Applicants' representatives can find no reference in this or any other section of Pettey to the local packetized interconnect that is tunneled to the receiving compute node being converted from a read response packet into a write request packet.

Therefore claim 33 is submitted to be further patentable over Pettey.

Claims 35 and 36

Claim 35 recites, inter alia, "passing the memory ID and the offset in the inter-node communication network packet to the receiving compute node". Claim 36 also incorporates this feature. The Office Action alleges that Pettey discloses this feature at col. 18, ln 55 - col. 19, ln. 36. The Applicants' representatives can find no reference in this or any other section of Pettey to

passing a memory ID and an offset calculated from an address of a local packetized interconnect packet to a receiving compute node.

Therefore claims 35 and 36 are submitted to be further patentable over Pettey.

Claim 38

Claim 38 recites, at the network interface of the receiving compute node, encapsulating a local packetized interconnect read request packet and forwarding the encapsulated read request packet to the sending compute node by way of the inter-node communication network. The Office Action alleges that Pettey discloses this feature at col. 18, ln 55 - col. 19, ln. 36. The Applicants' representatives can find no reference in this or any other section of Pettey to encapsulating a local packetized interconnect read request packet at a network interface of a receiving compute node and sending the encapsulated local packetized interconnect read request packet to a network interface of a sending compute node.

Therefore claim 38 is submitted to be further patentable over Pettey.

Claims 39-41 and 43-46

Claims 39, 43 and 45 are independent. These independent claims have been amended similarly to claim 1. These claims are submitted to be patentable over Pettey for the reasons set out above in relation to claim 1.

Compliance with 35 U.S.C. §103

The Office Action cites US patent 7,024,613 (Carnevale) in combination with Pettey in relation to claims 12 to 18. The Applicant submits that claims 12 to 18 are patentable over the cited combination.

Claims 12 to 18 depend from claim 1. Carnevale does not remedy the above-noted defects of Pettey. For at least this reason, claims 12 to 18 are not rendered obvious by the cited combination of Pettey with Carnevale.

Claims 12 to 18 are further submitted to distinguish Carnevale for the reasons set out below.

Claim 12

Claim 12 recites "generating a write confirmation packet at the network interface of the sending compute node and dispatching the write confirmation packet on the local packetized interconnect of the sending compute node". The Office Action indicates that this feature is disclosed in Carnevale at col. 4, ln. 1 - 24. The Applicant respectfully submits that this is incorrect.

The quoted material in Carnevale describes acknowledge packets being sent from an IB responder back to an IB requester (see col. 3, ln. 67 to col. 4, ln. 1). The IB responder is at a receiving end of the IB link. Carnevale, as understood, fails to disclose the feature of claim 12.

For this additional reason, Claim 12 is submitted to patentably distinguish the cited combination of Pettey and Carnevale.

Claims 14 to 16

Claim 14 recites "changing the write confirmation packet by altering the contained address of the memory location in the address space of the receiving compute node to a corresponding address in the first range of addresses". Carnevale, as understood, fails to disclose or suggest this feature. As noted above, Pettey also fails to disclose this feature.

For this additional reason, Claim 14, and claims 15-16 which depend from claim 14 are submitted to patentably distinguish the cited combination of Pettey and Carnevale.

Conclusion

The applicant submits that pending claims 1 to 3, 5 to 41 and 43 to 46, as amended, are in condition for allowance. Reconsideration and allowance of this application are respectfully requested.

Respectfully submitted,

By: /GavinNManning/
Gavin N. Manning
Registration No. 36,412
tel: 604.669.3432 ext. 9046
fax: 604.681.4081
e-mail: GNMDocket@patentable.com